arXiv:1510.06762v1 [cs.DM] 22 Oct 2015

Random-Cluster Dynamics in $\mathbb{Z}^2$

Antonio Blanca* Alistair Sinclair†

Abstract

The random-cluster model has been widely studied as a unifying framework for random graphs, spin systems and electrical networks, but its dynamics have so far largely resisted analysis. In this paper we analyze the Glauber dynamics of the random-cluster model in the canonical case where the underlying graph is an $n \times n$ box in the Cartesian lattice $\mathbb{Z}^2$. Our main result is a $O(n^2 \log n)$ upper bound for the mixing time at all values of the model parameter p except the critical point $p = p_c(q)$, and for all values of the second model parameter $q \geq 1$. We also provide a matching lower bound proving that our result is tight. Our analysis takes as its starting point the recent breakthrough by Beffara and Duminil-Copin on the location of the random-cluster phase transition in $\mathbb{Z}^2$. It is reminiscent of similar results for spin systems such as the Ising and Potts models, but requires the reworking of several standard tools in the context of the random-cluster model, which is not a spin system in the usual sense.

*Computer Science Division, U.C. Berkeley, Berkeley, CA 94720. Email: ablanca@cs.berkeley.edu. Research supported in part by an NSF Graduate Research Fellowship and NSF grant CCF-1420934.

†Computer Science Division, U.C. Berkeley, Berkeley, CA 94720. Email: sinclair@cs.berkeley.edu. Research supported in part by NSF grant CCF-1420934.

1 Introduction

Let $G = (V, E)$ be a finite graph. The *random-cluster model* on G with parameters $p \in (0, 1)$ and $q > 0$ assigns to each subgraph $(V, A \subseteq E)$ a probability

$$\mu_{G,p,q}(A) \propto p^{|A|}(1-p)^{|E|-|A|}q^{c(A)}, \tag{1}$$

where $c(A)$ is the number of connected components in (V, A). A is a *configuration* of the model.

The random-cluster model was introduced in the late 1960s by Fortuin and Kasteleyn [12] as a unifying framework for studying random graphs, spin systems in physics and electrical networks; see the book [15] for extensive background. When $q = 1$ this model corresponds to the standard Erdős-Rényi model on subgraphs of G, but when $q > 1$ (resp., $q < 1$) the probability measure favors subgraphs with more (resp., fewer) connected components, and is thus a strict generalization.

For the special case of integer $q \geq 2$ the random-cluster model is, in a precise sense, dual to the classical ferromagnetic *q-state Potts model*, where configurations are assignments of spin values $\{1, \ldots, q\}$ to the vertices of G; the duality is established via a coupling of the models (see, e.g., [10]). Consequently, the random-cluster model illuminates much of the physical theory of the Ising/Potts models. However, it should be stressed that the random-cluster model is not a "spin system" in the usual sense: in particular, the probability that an edge e belongs to A does not depend only on the dispositions of its neighboring edges but on the entire configuration A, since connectivity is a global property.

At the other extreme, when $q \to 0$, the set of (weak) limits that arise for various choices of p contains fundamental distributions on G, including uniform measures over the spanning trees, spanning forests and connected subgraphs of G.

The random-cluster model on $\mathbb{Z}^2$. The random-cluster model is well defined for the infinite 2-dimensional lattice graph $\mathbb{Z}^2$ as the limit of the sequence of random-cluster measures on $n \times n$ square regions Λ_n of $\mathbb{Z}^2$ as n goes to infinity. Recent breakthrough work of Beffara and Duminil-Copin [3] for the infinite measure has established the following *phase transition* at the critical value $p = p_c(q) = \sqrt{q}/(\sqrt{q}+1)$ for all $q \geq 1$: for $p < p_c(q)$ all connected components are finite with high probability; while for $p > p_c(q)$ there is at least one infinite component with high probability. It was also established in [3] that for $p < p_c(q)$ the model exhibits "decay of connectivities", i.e., the probability that two vertices lie in the same connected component decays to zero exponentially with the distance between them. This property is analogous to the classical "decay of correlations" that has long been known for the Ising model (see, e.g., [21]).

In this paper, we explore the consequences of the Beffara-Duminil-Copin result for the *dynamics* of the model in the case most widely studied in the literature, when G is an $n \times n$ square region Λ_n of $\mathbb{Z}^2$ and $q \geq 1$.

Glauber dynamics. A *Glauber dynamics* for the random-cluster model is any local Markov chain on configurations that is reversible w.r.t. the measure (1), and hence converges to it. Specifically we will consider the "heat-bath" dynamics, which at each step updates one edge of the current configuration A as follows:

(i) pick an edge $e \in E$ u.a.r.;

(ii) replace A by $A \cup \{e\}$ with probability

$$\frac{\mu_{G,p,q}(A \cup \{e\})}{\mu_{G,p,q}(A \cup \{e\}) + \mu_{G,p,q}(A \setminus \{e\})};$$

(iii) else replace A by $A \setminus \{e\}$.

These transition probabilities can be easily computed, as explained in Section 2.

Glauber dynamics for spin systems have been widely studied in both statistical physics and computer science. On the one hand, they provide a Markov chain Monte Carlo algorithm for sampling configurations of the system from the Gibbs distribution; on the other hand, they are a generally accepted model for the evolution of the underlying physical system. The primary object of study is the *mixing time*, i.e., the number of steps until the dynamics is close to its stationary distribution, starting from any initial configuration.

There has been much activity over the past two decades in analyzing Glauber dynamics for spin systems such as the Ising and Potts models, and deep connections have emerged between the mixing time and the phase structure of the physical model. In contrast, the Glauber dynamics for the random-cluster model remains very poorly understood. The main reason for this appears to be the fact mentioned above that connectivity is a global property; this has led to the lack (until the recent breakthrough [3]) of a precise understanding of the phase transition, as well as the failure of existing Markov chain analysis tools.

Essentially all existing bounds on the mixing time of the random-cluster Glauber dynamics (even in the simplest case of the mean-field model, where G is the complete graph [5]) are indirect, and proceed via a non-local dynamics (the so-called "Swendsen-Wang" dynamics [26] or its variants). Comparison technology developed recently by Ullrich [27–29] allows bounds for the Glauber dynamics of the Ising/Potts models to be translated to the Swendsen-Wang dynamics, and then again to the random-cluster dynamics. This leads, for example, to an upper bound of $O(n^6 \log^2 n)$ on the mixing time of the random-cluster dynamics in $\Lambda_n \subset \mathbb{Z}^2$, at all values $p \neq p_c(q)$ for all integer $q \geq 2$. This approach has several serious limitations:

1. The comparison method invokes linear algebra, and hence suffers an inherent penalty of at least $\Omega(n^4)$ in the mixing time bound; thus tight bounds can never be obtained in this way.

2. The comparison method also yields no insight into the actual behavior of the random-cluster dynamics, so, e.g., it is unlikely to illuminate the connections with phase transitions.

3. Since it relies on comparison with the Ising/Potts models, this analysis applies only for *integer* values of q, while the random-cluster model is defined for all positive values of q.

Results. In this paper we present the first direct analysis of the random-cluster dynamics and prove the following tight theorem for the important case of $\mathbb{Z}^2$:

Theorem 1.1. *For any $q \geq 1$, the mixing time of the Glauber dynamics for the random-cluster model on $\Lambda_n \subset \mathbb{Z}^2$ is $\Theta(n^2 \log n)$ at all values of $p \neq p_c(q)$.*

Theorem 1.1, as stated, holds for the random-cluster model with so-called "free" boundary conditions (i.e., there are no edges in $\mathbb{Z}^2 \setminus \Lambda_n$). In fact, as a consequence of our proof, it also holds for the case of "wired" boundary conditions (in which all vertices on the external face of Λ_n are connected).

The main component of our result is the analysis of the sub-critical regime $p < p_c$; the result for the super-critical regime $p > p_c$ follows from it easily by self-duality of $\mathbb{Z}^2$ and the fact that p_c is exactly the self-dual point [3]. Our sub-critical upper bound analysis makes crucial use of the exponential decay of connectivities for $p < p_c$ established recently by Beffara and Duminil-Copin [3], as discussed earlier. This analysis is reminiscent of similar results for spin systems (such as the Ising model), in which exponential decay of *correlations* has been shown to imply rapid

mixing [22]. However, since the random-cluster model exhibits decay of connectivities rather than decay of correlations, we need to rework the standard tools used in these contexts. In particular, we make three innovations.

First, the classical notion of "disagreement percolation" [4], which is used to bound the speed at which influence can propagate in $\mathbb{Z}^2$ under the dynamics, has to be non-trivially extended to take account of the fact that in the random-cluster model influence spreads not from vertex to vertex but from cluster to cluster. Second, we need to translate the decay of connectivities in the infinite volume $\mathbb{Z}^2$ (as proved in [3]) to a stronger "spatial mixing" property in finite volumes Λ_n, with suitable *boundary conditions* around the external face; in doing this we use the machinery developed in [2] by Alexander, but adapted to hold for arbitrary (not just integer) q and for a suitable class of boundary conditions that we call "side-homogeneous" (see Section 2 for a definition). Finally, while we follow standard recursive arguments in relating the mixing time in Λ_n to that in smaller regions $\Lambda_{n'}$ for $n' \ll n$, our approach differs in its sensitivity to the boundary conditions on the smaller regions: previous applications for spin systems have typically required rapid mixing to hold in $\Lambda_{n'}$ for *arbitrary* boundary conditions, while in our case we require it to hold only for side-homogeneous conditions. This aspect of our proof is actually essential because the random-cluster model does *not* exhibit spatial mixing for arbitrary boundary conditions (see Section 4); our definition of side-homogeneous conditions is motivated by the fact that they are both restricted enough to allow spatial mixing to hold, and general enough to make the recursion go through. Our lower bound proof uses technology from analogous results for spin systems of Hayes and Sinclair [16], again adapted to the random-cluster setting.

Our results leave open the question of the mixing time at the critical point $p = p_c(q)$. A full treatment of this regime may lie beyond the scope of current knowledge. In particular, the nature of the phase transition is not fully understood, but is conjectured to be second-order for $q \leq 4$ and first-order for $q > 4$; see [17] and [8] for partial results in this direction. This would suggest that the mixing time at p_c should be polynomial (though presumably of larger order than $\tilde{O}(n^2)$) for $q \leq 4$ and exponential in n for $q > 4$. Indeed, the former already follows for $q = 2$ from recent results of Lubetzky and Sly [20] on the Ising model at criticality, and the latter (for periodic boundary conditions) for the case of sufficiently large integer q from results of Borgs et al. [6,7] for the Glauber dynamics of the Potts model, using in both cases Ullrich's comparison techniques mentioned earlier.

The rest of the paper is organized as follows. We conclude this Introduction with a brief discussion of related work. Section 2 contains some basic terminology and facts used throughout the paper. In Sections 3 and 4 we derive our main analytical tools: bounds on the rate of disagreement percolation, and strong mixing for side-homogeneous boundary conditions, respectively. In Sections 5 and 6 we apply these tools to derive our mixing time bounds in the sub-critical regime $p < p_c(q)$. Finally, in Section 7 we use planar duality to extend the result to the super-critical regime $p > p_c(q)$, thus completing the proof of Theorem 1.1.

Related work. The random-cluster model has been the subject of extensive research in both the applied probability and statistical physics communities, which is summarized in the book by Grimmett [15]. A central open problem was to rigorously establish the phase transition in $\mathbb{Z}^2$ at $p_c(q) = \sqrt{q}/(\sqrt{q}+1)$, though this was not achieved until 2012 by Beffara and Duminil-Copin [3]. Predating the location of the phase transition, Alexander [2] showed that exponential decay of connectivities in $\mathbb{Z}^2$ (as was also established for $p < p_c(q)$ in [3]) implies spatial mixing in Λ_n for a certain class of boundary conditions and integer q.

The Glauber dynamics for the Ising model on $\mathbb{Z}^2$ is essentially completely understood thanks to decades of research. At all parameter values below the critical point β_c, the mixing time in Λ_n

is $O(n^2 \log n)$, while above β_c it is $\exp(\Omega(n))$ (see [21] for a comprehensive treatment, and also [20] for the behavior at β_c). Analogous results for the q-state Potts model follow from the random-cluster results of Beffara and Duminil-Copin [3] and Alexander [1], combined with the earlier work of Martinelli, Olivieri and Schonmann [22] relating spatial mixing to mixing times. Our work in the present paper carries through a parallel program for the random-cluster dynamics.

Glauber dynamics for the random-cluster model on any graph are poorly understood. As explained earlier, almost all the known results are derived by translating mixing time bounds from the Ising/Potts models using comparison techniques due to Ullrich [27–29]. Such translations typically incur a substantial overhead due to the use of linear algebra, hold only for integer q, and give no insight into the dynamics. One exception is a direct polynomial bound on the mixing time for the random-cluster dynamics on graphs with bounded tree-width due to Ge and Štefankovič [14].

Finally, we mention some results on *non-local* dynamics for the random-cluster model in the simpler *mean-field* case, where G is the complete graph K_n. The classical (non-local) Swendsen-Wang dynamics for the Ising and Potts models [26] may be viewed as a dynamics for the random-cluster model via the standard coupling [10], so recent tight mixing time results for this case [13,19] translate directly for integer values of q. The recent paper [5] proves similar results for the related Chayes-Machta dynamics, which applies to any (not necessarily integer) $q > 1$.

2 Preliminaries

In this section we gather a number of definitions and background results that we will refer to repeatedly. More details and proofs can be found in the books [15,18].

Random-cluster model on $\mathbb{Z}^2$. Let $\mathbb{L} = (\mathbb{Z}^2, \mathbb{E})$ be the square lattice graph, where for $u, v \in \mathbb{Z}^2$, $(u, v) \in \mathbb{E}$ iff $d(u, v) = 1$ with $d(\cdot, \cdot)$ denoting the Euclidean distance. Let $\Lambda_n \subseteq \mathbb{Z}^2$ be the set of vertices of $\mathbb{L}$ contained in a square box of side length n, and let $\Lambda = (\Lambda_n, E_n)$ be the graph whose edge set E_n contains all edges in $\mathbb{E}$ with both endpoints in Λ_n. We use $\partial\Lambda$ to denote the *boundary* of Λ; that is, the set of vertices in Λ_n connected by an edge in $\mathbb{E}$ to $\Lambda_n^c = \mathbb{Z}^2 \setminus \Lambda_n$.

A *random-cluster configuration* on Λ corresponds to a subset A of E_n. Alternatively, it is sometimes convenient to think of A as a vector in $\{0, 1\}^{|E_n|}$ indexed by the edges, where $A(e) = 1$ iff $e \in A$. Edges belonging to A are called *open*, and edges in $E_n \setminus A$ *closed*.

For any random-cluster configuration A^c on Λ_n^c, we may consider the conditional random-cluster measure induced in Λ_n by A^c. To make this precise, we introduce the standard concept of *boundary conditions*. A boundary condition for Λ is a partition $\eta = (P_1, P_2, ..., P_k)$ of $\partial\Lambda$ which encodes how the vertices of $\partial\Lambda$ are connected in a fixed configuration A^c on Λ_n^c; i.e., for all $u, v \in \partial\Lambda$, $u, v \in P_i$ iff u and v are connected by a path in A^c (see Figure 1(a)). In this case we also say that u and v are *wired* in η.

For $A \subseteq E_n$ and a boundary condition η, let $c(A, \eta)$ be the number of connected components of (Λ_n, A) when the connectivities from the boundary condition η are also considered. More precisely, if C_1, C_2 are connected components of A, and there exist $u \in C_1 \cap \partial\Lambda$ and $v \in C_2 \cap \partial\Lambda$ such that u and v are wired in η, then C_1 and C_2 are identified as the same connected component in A. The random-cluster measure on Λ with boundary condition η and parameters $p \in (0, 1)$ and $q > 0$ is then given by

$$\mu_{\Lambda,p,q}^{\eta}(A) = \frac{p^{|A|}(1-p)^{|E_n \setminus A|} q^{c(A,\eta)}}{Z_{\Lambda,p,q}^{\eta}}, \tag{2}$$

where $Z_{\Lambda,p,q}^{\eta}$ is the normalizing constant, or *partition function*. (Cf. equation (1) in the Introduction, which corresponds to the special case when the boundary condition η is "free"; see below.) When

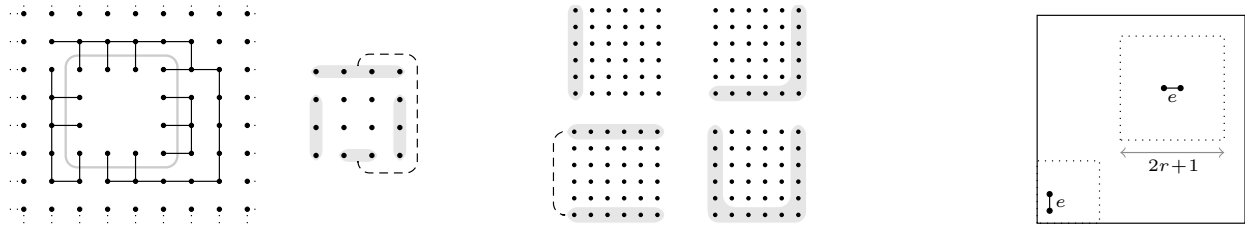

(a)(i) (a)(ii) (b) (c)

Figure 1: (a): (i) $\Lambda_4 \subset \mathbb{Z}^2$ *with a random-cluster configuration* A^c *in* Λ_4^c, (ii) *the boundary condition induced in* Λ_4 *by* A^c*;* (b) *examples of side-homogeneous boundary conditions;* (c) $B(e,r)$ *for two edges* e *of* Λ.

Λ, p and q are clear from the context we will just write μ^η.

Free, wired and side-homogeneous boundary conditions. Some boundary conditions will be of particular interest to us. In the *free* boundary condition no two vertices of $\partial\Lambda$ are wired. At the other extreme, in the *wired* boundary condition all vertices of $\partial\Lambda$ are pairwise wired. We will use $\mu^0_{\Lambda,p,q}$ and $\mu^1_{\Lambda,p,q}$ to denote the random-cluster measures on Λ with free and wired boundary conditions, respectively.

We also consider a class of boundary conditions which we call *side-homogeneous*. Let L_1, L_2, L_3, $L_4 \subset \partial\Lambda$ be the sets of vertices on each *side* of the square box Λ_n. (A corner vertex of Λ_n belongs to two sides.) The class of side-homogeneous boundary conditions contains all $\eta = (P_1, ..., P_k)$ satisfying:

(P1) $|P_i| > 1$ for at most one i; and

(P2) If $|P_i| > 1$, then P_i is the union of some of the sets L_j; i.e., $P_i = \bigcup_{j\in\kappa} L_j$, for some $\kappa \subseteq \{1,2,3,4\}$.

(See Figure 1(b).) Note that both the free and wired boundary conditions are side-homogeneous and there are in total 16 distinct side-homogeneous boundary conditions.

Monotonicity. For any pair of boundary conditions η and ψ, we say $\eta \le \psi$ if the partition η is a refinement of ψ; i.e., if the connectivities induced by η in $\partial\Lambda$ are also induced by ψ. When $q \ge 1$, $\eta \le \psi$ implies $\mu^\eta_{\Lambda,p,q} \preceq \mu^\psi_{\Lambda,p,q}$, where $\preceq$ denotes *stochastic domination*; i.e., $\mu^\eta_{\Lambda,p,q}(\mathcal{E}) \le \mu^\psi_{\Lambda,p,q}(\mathcal{E})$ for all increasing events $\mathcal{E}$. (An event $\mathcal{E}$ is *increasing* if it is preserved by the addition of edges.)

Planar duality. Let $\Lambda^* = (\Lambda_n^*, E_n^*)$ denote the planar dual of Λ in the usual sense. That is, Λ_n^* corresponds to the set of faces of Λ, and for each $e \in E_n$, there is a dual edge $e^* \in E_n^*$ connecting the two faces bordering e. The random-cluster measure satisfies $\mu_{\Lambda,p,q}(A) = \mu_{\Lambda^*,p^*,q}(A^*)$, where A^* is the dual configuration to A (i.e., $e^* \in E_n^*$ iff $e \in E_n$), and

$$p^* = \frac{q(1-p)}{p + q(1-p)}.$$

(This duality relation is a consequence of Euler's formula.) The unique value of p satisfying $p = p^*$, denoted $p_{sd}(q)$, is called the *self-dual point*.

Infinite measure and phase transition. The random-cluster measure may also be defined on the infinite lattice $\mathbb{Z}^2$ by considering the sequence of random-cluster measures on Λ_n with free boundary

conditions as $n \to \infty$. This sequence converges to a limiting measure $\mu_{\mathbb{L},p,q}$, which is known as the *random-cluster measure on* $\mathbb{L}$. The measure $\mu_{\mathbb{L},p,q}$ exhibits a phase transition corresponding to the appearance of an infinite connected component. That is, there exists a critical value $p = p_c(q)$ such that if $p < p_c(q)$ (resp., $p > p_c(q)$), then all components are finite (resp., there is at least one infinite component) w.h.p.

For $q \geq 1$, the exact value of $p_c(q)$ for $\mathbb{L}$ was only recently settled in breakthrough work by Beffara and Duminil-Copin [3], who proved the long standing conjecture

$$p_c(q) = p_{sd}(q) = \frac{\sqrt{q}}{\sqrt{q}+1}.$$

Exponential decay of connectivies and spatial mixing. In [3], it was also established that the phase transition is very sharp, meaning that as soon as $p < p_c(q)$ there is exponential decay of connectivities. More formally, for $q \geq 1$ and any fixed $p < p_c(q)$, there exist positive constants C, λ such that for all $u, v \in \mathbb{Z}^2$,

$$\mu_{\mathbb{L},p,q}(u \leftrightarrow v) \leq C\mathrm{e}^{-\lambda\, d(u,v)},$$

where $u \leftrightarrow v$ denotes the event that u and v are connected by a path of open edges. In work predating [3], Alexander [2] showed that exponential decay of connectivities implies exponential decay of *finite volume* connectivities uniformly over all boundary conditions. That is, for any boundary condition η on Λ, and all $u, v \in \Lambda_n$,

$$\mu^{\eta}_{\Lambda,p,q}(u \overset{\Lambda}{\leftrightarrow} v) \leq C\mathrm{e}^{-\lambda\, d(u,v)}, \tag{3}$$

where $u \overset{\Lambda}{\leftrightarrow} v$ is the event that u and v are connected by a path of open edges in Λ.

The notion of decay of connectivities for the random-cluster model is analogous to the notion of decay of correlations in spin systems, which is ubiquitous in the spin systems literature. As in spin systems, we require in our analysis of the dynamics a stronger form of decay of connectivities known as *spatial mixing*.

For $e \in E_n$, let $B(e, r) \subset \Lambda_n$ be the set of vertices in the minimal square box around e such that $d(\{e\}, v) \geq r$ for all $v \in \Lambda_n \setminus B(e, r)$. Note that if $d(\{e\}, \partial\Lambda) > r$, then $B(e, r)$ is just a square box of side length $2r + 1$ centered at e; otherwise $B(e, r)$ intersects $\partial\Lambda$ (see Figure 1(c)). Let $E(e, r)$ be the set of edges in E_n with both endpoints in $B(e, r)$, and let $E^c(e, r) = E_n \setminus E(e, r)$. The spatial mixing property we use, which is slightly weaker than that defined in [2], states that for all $e \in E_n$ and for every pair of configurations A^c_1, A^c_2 on $E^c = E^c(e, r)$, we have

$$\left|\mu^{\eta}_{\Lambda,p,q}(\, e = 1 \mid A^c_1\,) - \mu^{\eta}_{\Lambda,p,q}(\, e = 1 \mid A^c_2\,)\right| \leq \mathrm{e}^{-\lambda r} \tag{4}$$

for some constant $\lambda > 0$. Alexander [2] showed that (3) implies (4) for a certain class of boundary conditions η when q is an integer. In Section 4 we will show, using the machinery developed in [2], that (4) holds for all side-homogeneous boundary conditions η for any (not necessarily integer) $q \geq 1$.

Glauber dynamics. The *Glauber dynamics*, which we denote $\mathcal{M}$, is a local Markov chain on the random-cluster configurations of $\Lambda = (\Lambda_n, E_n)$ that is reversible with respect to $\mu^{\eta}_{\Lambda,p,q}$ for any boundary condition η; as a result, its stationary distribution is $\mu^{\eta}_{\Lambda,p,q}$. Given a random-cluster configuration $A_t \subseteq E_n$ at time t, a step of $\mathcal{M}$ results in a new configuration A_{t+1} as follows:

(i) pick $e \in E_n$ u.a.r;

(ii) let $A_{t+1} = A_t \cup \{e\}$ with probability

$$\begin{cases} \frac{p}{p+q(1-p)} & \text{if } e \text{ is a "cut edge" in } (\Lambda_n, A_t); \\ p & \text{otherwise;} \end{cases}$$

(iii) else let $A_{t+1} = A_t \setminus \{e\}$.

(We say e is a *cut edge* in (Λ_n, A_t) iff changing the current configuration of e changes the number of connected components of A_t.) Note that this definition of the Glauber dynamics is equivalent to that in the Introduction, as can be seen by computing the ratios of the appropriate stationary weights $\mu^{\eta}_{\Lambda,p,q}(\cdot)$.

Mixing time and couplings. Let Ω_{RC} be the set of random-cluster configurations of Λ, and let $\mathcal{M}^t(X_0, \cdot)$ be the distribution of $\mathcal{M}$ after t steps starting from $X_0 \in \Omega_{\text{RC}}$. Let

$$\tau_{\text{mix}}(\varepsilon) := \max_{X_0 \in \Omega_{\text{RC}}} \min_t \left\{ ||\mathcal{M}^t(X_0, \cdot) - \mu^{\eta}(\cdot)||_{\text{TV}} \leq \varepsilon \right\},$$

where $||\cdot||_{\text{TV}}$ denotes total variation distance. The *mixing time* of $\mathcal{M}$ is given by $\tau_{\text{mix}} := \tau_{\text{mix}}(1/4)$. It is well-known that $\tau_{\text{mix}}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil \, \tau_{\text{mix}}$ for any positive $\varepsilon < 1/2$.

A *(one step) coupling* of the Markov chain $\mathcal{M}$ specifies, for every pair of states $(X_t, Y_t) \in \Omega^2_{\text{RC}}$, a probability distribution over (X_{t+1}, Y_{t+1}) such that the processes $\{X_t\}$ and $\{Y_t\}$, viewed in isolation, are faithful copies of $\mathcal{M}$, and if $X_t = Y_t$ then $X_{t+1} = Y_{t+1}$. The *coupling time*, denoted T_{coup}, is the minimum T such that $\Pr[X_T \neq Y_T] \leq 1/4$, starting from the worst possible pair of configurations X_0, Y_0. The following inequality is standard : $\tau_{\text{mix}} \leq T_{\text{coup}}$ (see, e.g., [18]).

One coupling will be of particular interest to us. Namely, consider the coupling that couples the evolution of two copies of $\mathcal{M}$, $\{X_t\}$ and $\{Y_t\}$, by using the same random $e \in E_n$ in step (i) and the same uniform random number $r \in [0, 1]$ to decide whether to add or remove e in steps (ii) and (iii). We call this the *identity coupling*. It is straightforward to verify that, when $q \geq 1$, the identity coupling is a *monotone coupling*, in the sense that if $X_t \subseteq Y_t$ then $X_{t+1} \subseteq Y_{t+1}$ with probability 1. In fact, the identity coupling can be extended to a simultaneous coupling of *all* configurations that preserves the partial order $\subseteq$. Therefore, the coupling time starting from any pair of configurations is bounded by the coupling time for initial configurations $Y_0 = \emptyset$ and $X_0 = E_n$, which are the unique minimal and maximal elements in the partial order [24].

3 The speed of disagreement percolation

In spin systems, a central idea in the analysis of local Markov chains is to bound the speed at which disagreements propagate. Bounds of this sort can be deduced from *disagreement percolation* (or *path of disagreements*) arguments (see, e.g., [4,9]). These arguments are based on the idea that in spin systems interactions only occur between neighboring sites, and thus if two configurations agree everywhere except in some region A, then it takes many steps for a local Markov chain under the identity coupling to propagate these disagreements to regions that are far from A.

In this section, we provide a bound on the speed of propagation of disagreements for the Glauber dynamics of the random-cluster model on $\Lambda = (\Lambda_n, E_n)$, under side-homogeneous boundary conditions. The random-cluster model has long range interactions in the form of arbitrarily long paths, so disagreements could potentially propagate arbitrarily fast. Our insight is to restrict attention to pairs of configurations where one of the configurations is *stationary*; i.e., it has law $\mu^{\eta}_{\Lambda,p,q}$. Then by (3), the probability of long paths decays exponentially with the length of the path, which makes the long range interactions manageable.

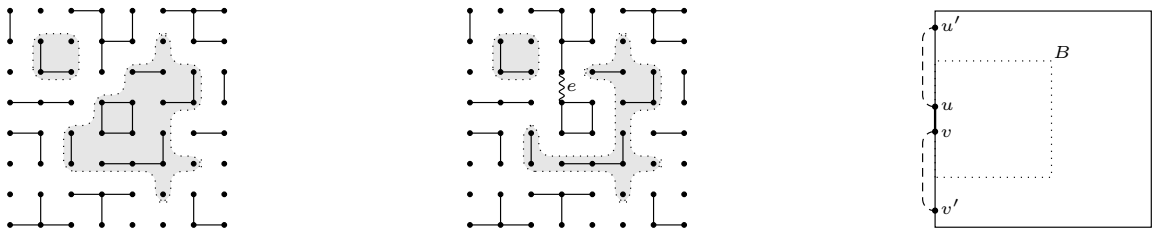

(a) (b) (c)

Figure 2: (a) *a random-cluster configuration A in B, with $R_0 = \Gamma(A, B)$ (shaded);* (b) R_1*, assuming edge e was updated in the first step;* (c) *a boundary condition ψ where the spatial mixing property does not hold.*

We will use the notation introduced Section 2. Also, for a random-cluster configuration A on Λ and $D \subseteq \Lambda_n$, we use $A(D)$ to denote the configuration induced by A on the edges with both endpoints in D.

Lemma 3.1. *Let $p < p_c(q)$, $q \geq 1$ and consider two copies $\{X_t\}$, $\{Y_t\}$ of the Glauber dynamics in $\Lambda = (\Lambda_n, E_n)$ with a side-homogeneous boundary condition η. Assume Y_0 has law $\mu^{\eta}_{\Lambda,p,q}$ and that $X_0(B(e,r)) = Y_0(B(e,r))$ for some $e \in E_n$ and $r \geq 1$. If the evolutions of $\{X_t\}$ and $\{Y_t\}$ are coupled using the identity coupling, then there exist absolute constants $c, C, \lambda > 0$ such that, for $m = |E_n|$, $r \geq c$ and $1 \leq k \leq r^{1/4}/(4\mathrm{e}^2)$, we have*

$$\Pr[\, X_{km}(e) \neq Y_{km}(e) \,] \;\leq\; C \exp(-\lambda r^{1/4}).$$

Proof. Let ∂B denote the boundary of $B = B(e, r)$; that is, ∂B is the set of vertices in B connected by an edge in E_n to $\Lambda_n \setminus B$. (Note that if $v \in B \cap \partial\Lambda$, then $v \notin \partial B$.) For some fixed ℓ to be chosen later, and each $t \geq 0$, consider the event

$$\mathcal{E}_{\ell,t} := \{u \overset{Y_t}{\nleftrightarrow} v \quad \forall u, v \in B \text{ s.t. } d(u, v) > \ell\},$$

where $u \overset{Y_t}{\nleftrightarrow} v$ denotes the event that u and v are not connected by a path in $Y_t(B)$. Let $\mathcal{E}_\ell := \bigcap_{t=0}^{km} \mathcal{E}_{\ell,t}$. Then,

$$\Pr[\, X_{km}(e) \neq Y_{km}(e) \,] \leq \Pr[\, X_{km}(e) \neq Y_{km}(e) \,|\, \mathcal{E}_\ell \,] + \Pr[\, \neg\mathcal{E}_\ell \,]. \tag{5}$$

We bound each term on the right hand side of (5) separately.

First we introduce some notation. For any random-cluster configuration A of Λ, let

$$\Gamma(A, B) := B \setminus \bigcup\nolimits_{v \in \partial B} C(v, A), \tag{6}$$

where $C(v, A)$ is the set of vertices in the connected component of v in (Λ_n, A). Also, for $D \subseteq \Lambda_n$, let $\partial E(D) := \{(u, v) \in E_n : u \in D, v \notin D\}$.

Consider the sequence of regions $R_0 \supseteq R_1 \supseteq \ldots$, such that $R_0 = \Gamma(Y_0, B)$ and

$$R_{t+1} = \begin{cases} R_t & \text{if no edge from } \partial E(R_t) \text{ is updated at time } t; \\ \\ R_t \setminus C(a_t, Y_t) & \text{if } (a_t, b_t) \in \partial E(R_t) \text{ with } a_t \in R_t,\, b_t \notin R_t \text{ is the edge updated at time } t; \end{cases}$$

(see Figures 2(a), 2(b)). Note that R_t need not be a connected region of $\mathbb{Z}^2$ and that the second case above applies regardless of whether $(a_t, b_t) \in Y_{t+1}$.

The key observation is that for all $t \geq 0$, R_t is a region of B in which $X_t(R_t) = Y_t(R_t)$. To see this, by induction on t assume $X_t(R_t) = Y_t(R_t)$ and let ψ_{X} (resp., ψ_{Y}) be the boundary condition induced in R_t by $X_t(\Lambda_n \setminus R_t)$ (resp., $Y_t(\Lambda_n \setminus R_t)$) and η. (Observe that ψ_{X} and ψ_{Y} are partitions of $\partial R_t \cup (R_t \cap \partial\Lambda)$.) We consider three cases based on the edge (a_t, b_t) updated at time t. First, if $\{a_t, b_t\} \cap R_t = \emptyset$, then clearly $X_{t+1}(R_{t+1}) = Y_{t+1}(R_{t+1})$. The second possibility is that $\{a_t, b_t\} \cap R_t = \{a_t\}$; if this is the case, then $R_{t+1} = R_t \setminus C(a_t, Y_t)$, $C(a_t, X_t) = C(a_t, Y_t)$ and as a result $X_{t+1}(R_{t+1}) = Y_{t+1}(R_{t+1})$.

Finally, when $\{a_t, b_t\} \subseteq R_t$ we show that $\psi_{\mathrm{X}} = \psi_{\mathrm{Y}}$, from which it follows that $X_{t+1}(R_{t+1}) = Y_{t+1}(R_{t+1})$. First observe that every edge in $\partial E(R_t)$ is closed in both X_t and Y_t. Thus, if $|R_t \cap \partial\Lambda| \leq 1$, then ψ_{X}, ψ_{Y} are both the free boundary condition on R_t and we are done. Hence, assume that $|R_t \cap \partial\Lambda| \geq 2$ and let u, v be any two distinct vertices in $R_t \cap \partial\Lambda$. If u and v are wired in η, then they are also wired in ψ_{X}. Moreover, if u and v are not wired in η, then property (P1) implies that one of them (say, v) is necessarily a singleton element in η. Thus, v is also a singleton element in ψ_{X}, since there are no paths from v to $\Lambda_n \setminus R_t$ in X_t. Hence, u and v are wired in ψ_{X} (resp., ψ_{Y}) iff they are wired in η. Since also every vertex in $\partial R_t \setminus \partial\Lambda$ is a singleton in both ψ_{X} and ψ_{Y}, we conclude that $\psi_{\mathrm{X}} = \psi_{\mathrm{Y}}$.

Since $X_t(R_t) = Y_t(R_t)$ for all $t \geq 0$, if both endpoints of e lie in R_{km} then $X_{km}(e) = Y_{km}(e)$. Also, since we are conditioning on $\mathcal{E}_\ell$ and we will choose $\ell \ll r$, both endpoints of e lie in R_0. So, if $X_{km}(e) \neq Y_{km}(e)$, we may take v_0 to be the first endpoint of e to be removed from R_{t_0}, at some time $t_0 \leq km$. Let $e_1 \in \partial E(R_{t_0})$ be the unique edge whose update is responsible for removing v_0 from R_{t_0}. Starting from e_1 we can then construct a sequence of edges $e_1, e_2, \ldots$ such that $e_i = (u_i, v_i)$, with $u_i \in R_{t_{i-1}}$ and $v_i \notin R_{t_{i-1}}$, is the edge that removes v_{i-1} from $R_{t_{i-1}}$ at time t_{i-1}. Note that $t_0 > t_1 > \ldots$ and that the sequence $e_1, e_2, \ldots$ stops once it reaches an edge $e_{\tilde{t}}$ incident to a vertex outside R_0. We call the sequence $e_1, e_2, \ldots, e_{\tilde{t}}$ a *witness* for the fact that $X_{km}(e) \neq Y_{km}(e)$.

The vertices v_i and u_{i+1} are in the same connected component at time t_i, and since we are conditioning on $\mathcal{E}_\ell$, we have $d(v_i, u_{i+1}) \leq \ell$. Therefore, the number of witnesses of length L is (crudely) at most $(4(\ell+1)^2)^L$. Note also that every witness contains at least $r/(\ell+1)$ edges, otherwise it cannot reach any of the vertices outside of R_0. Moreover, the probability that a given witness of length L is updated by the identity coupling in km steps is $\binom{km}{L}\left(\frac{1}{m}\right)^L$. Hence,

$$\begin{aligned}
\Pr[\, X_{km}(e) \neq Y_{km}(e) \,|\, \mathcal{E}_\ell \,] \leq & \sum_{L \geq \frac{r}{\ell+1}} \binom{km}{L} \left(\frac{1}{m}\right)^L (4(\ell+1)^2)^L \\
\leq & \sum_{L \geq \frac{r}{\ell+1}} \left(\frac{4\mathrm{e}k(\ell+1)^2}{L}\right)^L \\
\leq & \ \omega^{\frac{r}{\ell+1}} \sum_{L \geq 0} \omega^L,
\end{aligned}$$

where $\omega = \frac{4\mathrm{e}k(\ell+1)^3}{r}$. By taking $\ell = r^{1/4} - 1$ and using the fact that $k \leq r^{1/4}/(4\mathrm{e}^2)$, we have

$$\Pr[\, X_{km}(e) \neq Y_{km}(e) \,|\, \mathcal{E}_\ell \,] \leq \frac{\mathrm{e}}{\mathrm{e}-1} \cdot \mathrm{e}^{-r^{3/4}}. \tag{7}$$

Now we turn our attention to the second term on the right hand side of (5). Let N be the number of updates the identity coupling performs in B up to time km, and let M be the number of

edges in B; i.e., $M := |E(e,r)| = \Theta(r^2)$. A Chernoff bound implies that $N > 2kM$ with probability $\exp(-\Omega(kM))$, and thus

$$\Pr[\,\neg\mathcal{E}_\ell\,] \le \Pr[\,\neg\mathcal{E}_\ell \,|\, N \le 2kM\,] + \mathrm{e}^{-\Omega(kM)}.$$

Observe that $\neg\mathcal{E}_\ell = \bigcup_{t=0}^{km} \neg\mathcal{E}_{\ell,t}$, and if the edge update at time t occurs outside B, we have $\neg\mathcal{E}_{\ell,t} = \neg\mathcal{E}_{\ell,t+1}$. Hence, a union bound implies

$$\Pr[\,\neg\mathcal{E}_\ell\,] \;\le\; 2kM \max_{0\le t\le km} \Pr[\,\neg\mathcal{E}_{\ell,t}\,] + \mathrm{e}^{-\Omega(kM)}. \tag{8}$$

To bound $\Pr[\,\neg\mathcal{E}_{\ell,t}\,]$ we use the exponential decay of finite volume connectivities (3). To do this, recall that Y_0 (and thus Y_t for all $t \ge 0$) has law μ^η. Also, there are only $O(r^4)$ pairs of vertices in B. Hence, (3) and a union bound imply

$$\Pr[\,\neg\mathcal{E}_{\ell,t}\,] \le O(r^4)\cdot \mathrm{e}^{-\Omega(\ell)} \le O(r^4)\cdot \mathrm{e}^{-\Omega(r^{1/4})}.$$

Since $1 \le k \le r^{1/4}/(4\mathrm{e}^2)$ and $M = \Theta(r^2)$, (8) gives

$$\Pr[\,\neg\mathcal{E}_\ell\,] \le O(r^{6.25})\cdot \mathrm{e}^{-\Omega(r^{1/4})} + \mathrm{e}^{-\Omega(r^2)}.$$

Together with (7), this implies that there exist constants $c, C, \lambda > 0$ such that for all $r \ge c$, we have

$$\Pr[\,X_{km}(e) \ne Y_{km}(e)\,] \le C\mathrm{e}^{-\lambda r^{1/4}},$$

as desired. □

4 Spatial mixing for side-homogeneous boundary conditions

In this section we show that the spatial mixing property (4) holds for the class of side-homogeneous boundary conditions on $\Lambda = (\Lambda_n, E_n)$. Let $e \in E_n$ and let $B = B(e,r)$ for some $r \ge 1$. Spatial mixing holds when the influence on e of the configuration in $E^c = E^c(e,r)$ decays exponentially with r. This is easy to establish when $B \cap \partial\Lambda = \emptyset$, since such influence is present only if there are paths of open edges from e to ∂B, and, by (3), the probability of such paths decays exponentially with r. However, if $B \cap \partial\Lambda \ne \emptyset$, the influence from E^c could also propagate along $B \cap \partial\Lambda$ via the boundary condition on Λ. This is why (4) does not hold for arbitrary boundary conditions, as the following concrete example illustrates.

With a slight abuse of notation, we use $\{E^c = 1\}$ (resp., $\{E^c = 0\}$) to denote the event that all the edges in E^c are open (resp., closed). Suppose $e = (u,v)$ is an edge in $\partial\Lambda$ that is far from the corners of Λ, and let ψ be the boundary condition on Λ where u is wired to a vertex $u' \in \partial\Lambda \setminus B$ and v is wired to a different vertex $v' \in \partial\Lambda \setminus B$ (see Figure 2(c)). When $p = 1/2$ and $q = 3$, we have $\mu^\psi(e = 1 | E^c = 1) = 1/2$. Also, by considering a small box around e, is easy to check that $\mu^\psi(e = 1 | E^c = 0) \le 2/5$. Both these bounds are independent of r and n; consequently, ψ does not have the spatial mixing property.

It turns out that side-homogeneous boundary conditions (and in particular property (P1)) rule out the possibility of influence propagating along $B \cap \partial\Lambda$. As a result, we are able to establish the spatial mixing property for side-homogeneous boundary conditions, as stated in the following lemma; the proof uses the machinery developed in [2].

Lemma 4.1. *Let $p < p_c(q)$, $q \geq 1$ and let η be a side-homogeneous boundary condition for $\Lambda = (\Lambda_n, E_n)$. For any $e \in E_n$, there exist constants $c, \lambda > 0$ such that for all $r \geq c$ and every pair of configurations A_1^c, A_2^c on E^c:*

$$|\mu^\eta(\, e = 1 \,|\, A_1^c \,) - \mu^\eta(\, e = 1 \,|\, A_2^c \,)| \;\leq\; \mathrm{e}^{-\lambda r}. \tag{9}$$

Proof. Consider the measure $\mu^w := \mu^\eta(\,\cdot\,|E^c = 1)$ on $B = B(e, r)$. Let $\Gamma(\cdot, \cdot)$ be defined as in (6). We derive the result from the following key fact, which we prove later.

Claim 4.2. *There exists a coupling π of $\mu^\eta(\,\cdot\,|\,A_1^c\,)$, $\mu^\eta(\,\cdot\,|\,A_2^c\,)$ and μ^w such that $\pi(A_1, A_2, A_w) > 0$ only if $A_1 \subseteq A_w$, $A_2 \subseteq A_w$ and A_1, A_2 agree on all edges with both endpoints in $\Gamma(A_w, B)$.*

Let $\Gamma^c(A_w, B) := \Lambda_n \setminus \Gamma(A_w, B)$. Given the coupling π, we have

$$\begin{aligned} |\mu^\eta(\, e = 1 \,|\, A_1^c \,) - \mu^\eta(\, e = 1 \,|\, A_2^c \,)| &\leq \pi(\, A_1(e) \neq A_2(e) \,) \\ &\leq \pi\,(\, \Gamma^c(A_w, B) \cap \{e\} \neq \emptyset \,) \\ &\leq \mu^w\,(\, \{e\} \leftrightarrow \partial B \,)\,, \end{aligned}$$

where $\{e\} \leftrightarrow \partial B$ denotes the event that there is a path from e to ∂B.

Now, exponential decay of finite volume connectivities (3) and a union bound over the boundary vertices imply

$$\mu^w(\, \{e\} \leftrightarrow \partial B \,) \leq 2C|\,\partial B|\mathrm{e}^{-\,\lambda\, r}.$$

Since $|\,\partial B| = \Theta(r)$, we obtain (9) for all $r \geq c$ for some constant $c > 0$, and hence the lemma. □

We conclude this section by providing the missing proof of Claim 4.2.

Proof. Let θ_1 (resp., θ_2) be the boundary condition induced on $B = B(e, r)$ by A_1^c (resp., A_2^c) and η. Note that μ^{θ_1}, μ^{θ_2} and μ^w are random-cluster measures on B with different boundary conditions, and clearly $\mu^w \succeq \mu^{\theta_1}$ and $\mu^w \succeq \mu^{\theta_2}$. Strassen's theorem (see, e.g., [25]) then implies the existence of monotone couplings μ_1 for μ^{θ_1} and μ^w, and μ_2 for μ^{θ_2} and μ^w. (Recall that μ_1 is a monotone coupling for μ^{θ_1} and μ^w iff every sample (A_{θ_1}, A_w) from μ_1 satisfies $A_{\theta_1} \subseteq A_w$.) We show next how to use μ_1 and μ_2 to construct the desired coupling π.

First, let $\Delta := \Gamma(A_w, B)$ and let ξ be the boundary condition induced in Δ by η and the configuration of A_w in $\Gamma^c(A_w, B)$. We construct π as follows:

(i) sample (A_{θ_1}, A_w) from μ_1;

(ii) sample A_{θ_2} from $\mu_2(\,\cdot\,|\,A_w\,)$; and

(iii) sample A_γ from $\mu^\xi_{\Delta,p,q}$.

Let π be the distribution of

$$(\{A_{\theta_1} \setminus E(\Delta)\} \cup A_\gamma, \{A_{\theta_2} \setminus E(\Delta)\} \cup A_\gamma, \{A_w \setminus E(\Delta)\} \cup A_\gamma)$$

after step (iii), where $E(\Delta)$ denotes the set of edges with both endpoints in Δ.

A straightforward calculation reveals that A_{θ_2} has law μ^{θ_2}, and thus after step (ii) the distribution of $(A_{\theta_1}, A_{\theta_2}, A_w)$ has all the right marginals. Moreover, since μ_1 and μ_2 are monotone couplings, $A_{\theta_1} \subseteq A_w$ and $A_{\theta_2} \subseteq A_w$.

We argue next that replacing the configuration in Δ with A_γ in step (iii) has no effect on the distribution. For this, let ξ_1 (resp., ξ_2) be the boundary condition induced in Δ by η and the configuration of A_{θ_1} (resp., A_{θ_2}) in $\Gamma^c(A_w, B)$. (Note that ξ, ξ_1, ξ_2 are partitions of $\partial\Delta \cup (\Delta \cap \partial\Lambda)$.)

We show that $\xi = \xi_1 = \xi_2$. This is easy to see when $\Delta \cap \partial\Lambda = \emptyset$, since in this case all three of them are the free boundary condition on Δ. This is because $A_{\theta_1} \subseteq A_w$, $A_{\theta_2} \subseteq A_w$ and every edge from $\partial\Delta$ to Δ^c is closed in A_w.

When $\Delta \cap \partial\Lambda$ is not trivial, only vertices in $\Delta \cap \partial\Lambda$ may be wired in ξ. Property (P1), together with the fact that every edge from $\partial\Delta$ to Δ^c is closed in A_w, implies that two vertices from $\Delta \cap \partial\Lambda$ are wired in ξ iff they are wired in η. The same holds for ξ_1 and ξ_2; hence, $\xi = \xi_1 = \xi_2$. (Note that this argument is essentially the same as the one used in the proof of Lemma 3.1 to show that the boundary conditions induced in R_t by X_t and Y_t are the same.)

Finally, the *domain Markov property* of random-cluster measures (see, e.g., [15]) ensures that indeed replacing Δ with A_γ has no effect on the distribution. Hence, π is a coupling of the measures μ^{θ_1}, μ^{θ_2}, and μ^w with all the desired properties. □

5 Mixing time upper bound in the sub-critical regime

In this section we prove our main result: the upper bound for the mixing time in Theorem 1.1 for $p < p_c(q)$. We state two theorems whose combination establishes the desired upper bound for $p < p_c(q)$. In Theorem 5.1 we show that spatial mixing for the class of side-homogeneous boundary conditions, as established in Section 4, implies a bound of $O(n^2 \log n (\log\log n)^2)$ for the mixing time of the Glauber dynamics on $\Lambda = (\Lambda_n, E_n)$, for any n and any side-homogeneous boundary condition. The proof is inductive and makes crucial use of property (P2) of side-homogeneous boundary conditions, which ensures that for any $e \in E_n$ and $r \geq 1$, the boundary condition induced in $B(e, r)$ by the events $\{E^c = 1\}$ or $\{E^c = 0\}$ is also side-homogeneous.

In Theorem 5.2 we show that a sufficiently good upper bound—in fact $O\left(n^{2.25}/\log n\right)$ suffices—on the mixing time of the Glauber dynamics can be bootstrapped to the desired upper bound of $O(n^2 \log n)$. The proof of Theorem 5.2 crucially uses the bounds on the speed of propagation of disagreements from Section 3. Our proofs are inspired by those in the spin systems literature, in particular those in [9,21,23].

Theorem 5.1. *Let $p < p_c(q)$, $q \geq 1$ and let η be a side-homogeneous boundary condition for $\Lambda = (\Lambda_n, E_n)$. There exists a fixed constant $C > 0$ such that for all n, the mixing time of the Glauber dynamics in Λ is at most $T(m) = Cm \log m (\log\log m)^2$, where $m := |E_n| = \Theta(n^2)$.*

Proof. We bound the coupling time T_{coup} of the identity coupling; the result then follows from the fact that $\tau_{\text{mix}} \leq T_{\text{coup}}$. Consider two copies $\{X_t\}$, $\{Y_t\}$ of the Glauber dynamics coupled with the identity coupling. We may assume $X_0 = E_n$ and $Y_0 = \emptyset$, since we know from Section 2 that this is the worst pair of starting configurations. We prove that

$$\Pr[\, X_T(e) \neq Y_T(e) \,] = O\left(m^{-2}\right)$$

for $T = T(m)$ and for all $e \in E_n$. The bound for the coupling time then follows from a union bound over the edges.

To bound $\Pr[\, X_T(e) \neq Y_T(e) \,]$, we introduce two additional copies $\{Z_t^+\}$, $\{Z_t^-\}$ of the Glauber dynamics. These two copies will only update edges with both endpoints in the box $B = B(e, r)$, for a suitable r we choose later. We set $Z_0^+ = X_0 = E_n$ and $Z_0^- = Y_0 = \emptyset$. The four Markov chains $\{X_t\}$, $\{Y_t\}$, $\{Z_t^+\}$ and $\{Z_t^-\}$ are coupled with the identity coupling, and the updates outside B are ignored by $\{Z_t^+\}$ and $\{Z_t^-\}$. By monotonicity of the identity coupling, we have $Z_t^- \subseteq Y_t \subseteq X_t \subseteq Z_t^+$ for all $t \geq 0$. Hence,

$$\begin{aligned}\Pr[\, X_t(e) \neq Y_t(e) \,] &\leq \Pr[\, Z_t^+(e) \neq Z_t^-(e) \,] \\ &= \Pr[\, Z_t^+(e) = 1 \,] - \Pr[\, Z_t^-(e) = 1 \,].\end{aligned}$$

The stationary measures of $\{Z_t^+\}$ and $\{Z_t^-\}$ are $\mu_\Lambda^\eta(\,\cdot\,|\,E^c=1\,)$ and $\mu_\Lambda^\eta(\,\cdot\,|\,E^c=0\,)$, respectively, where as usual $\{E^c=1\}$ (resp., $\{E^c=0\}$) denotes the event that all the edges in $E_n\setminus E(e,r)$ are open (resp., closed). The triangle inequality then implies

$$\begin{aligned}\Pr[\,X_t(e)\neq Y_t(e)\,] \;\le\; & \big|\Pr[\,Z_t^+(e)=1\,]-\mu_\Lambda^\eta(\,e=1\,|\,E^c=1\,)\big| && (10)\\ & +\big|\mu_\Lambda^\eta(\,e=1\,|\,E^c=1\,)-\mu_\Lambda^\eta(\,e=1\,|\,E^c=0\,)\big| && (11)\\ & +\big|\mu_\Lambda^\eta(\,e=1\,|\,E^c=0\,)-\Pr[\,Z_t^-(e)=1\,]\big|\,. && (12)\end{aligned}$$

The chains $\{Z_t^+\}$ and $\{Z_t^-\}$ are Glauber dynamics on the smaller square box B. Also, the boundary conditions induced in B by η and the events $\{E^c=1\}$, $\{E^c=0\}$ are side-homogeneous. Hence, we proceed inductively.

First note that for any fixed m_0, the result holds for all square boxes of volume at most m_0 by simply adjusting the constant $C=C(m_0)$. Now, let $r=c\log m$ for some constant c we choose later, and assume the mixing time bound holds for square boxes of volume $M:=|E(e,r)|$ with side-homogeneous boundary condition. After $T(m)$ steps, the expected number of updates in B is

$$T(m)\frac{M}{m}=CM\log m(\log\log m)^2\ge\lceil 4\log_2 m\rceil T(M),$$

where we choose m_0 such that the last inequality holds for all $m>m_0$. Hence, a Chernoff bound implies that the number of updates in B is at least $\lceil 2\log_2 m\rceil T(M)$ with probability at least $1-m^{-2}$.

The induction hypothesis implies that the mixing time of $\{Z_t^+\}$ is at most $T(M)$. Hence, if indeed $\lceil 2\log_2 m\rceil T(M)$ updates happen in B, then

$$||\,Z_t^+(\cdot)-\mu_\Lambda^\eta(\,\cdot\,|\,E^c=1\,)\,||_{\mathrm{TV}}\le\frac{1}{m^2}.$$

(Here we used the fact from Section 2 that $\tau_{\mathrm{mix}}(\varepsilon)\le\lceil\log_2\varepsilon^{-1}\rceil\,\tau_{\mathrm{mix}}$.) Combining this with the above Chernoff bound, we have

$$\big|\Pr[\,Z_t^+(e)=1\,]-\mu_\Lambda^\eta(\,e=1\,|\,E^c=1)\big|\le\frac{2}{m^2}.$$

The quantity in (12) is bounded similarly.

Finally, taking $c=2/\lambda$, the spatial mixing property (Lemma 4.1) implies that (11) is at most $1/m^2$. Putting these bounds together we get

$$\Pr[\,X_t(e)\neq Y_t(e)\,]\le\frac{5}{m^2},$$

as desired. □

Theorem 5.2. *Let $p<p_c(q)$, $q\ge1$ and let m_0, c be sufficiently large and sufficiently small constants, respectively. Assume that the mixing time of the Glauber dynamics in any square box of volume m_0 with side-homogeneous boundary conditions is at most $\frac{c\,m_0^{9/8}}{\log_2 m_0}$. Then, the mixing time of the Glauber dynamics in Λ with side-homogeneous boundary conditions is $O(n^2\log n)$.*

Proof. Let $m:=|E_n|=\Theta(n^2)$ and the let η be a side-homogeneous boundary condition for Λ. Also, let $\{X_t\}$, $\{Y_t\}$ be two copies of the Glauber dynamics in Λ coupled with the identity coupling. We prove that for $1\le k=o(m^{1/8})$, we have

$$\Pr[X_{km}(e)\neq Y_{km}(e)]\le\mathrm{e}^{-\Omega(k)} \qquad (13)$$

for any $e \in E_n$ and any pair of initial configurations X_0, Y_0. Hence, $\Pr[X_{km}(e) \neq Y_{km}(e)] \leq 1/(4m)$ for some $k = O(\log m)$, and a union bound over the edges implies $\tau_{\mathrm{mix}} \leq T_{\mathrm{coup}} = O(m \log m)$, as desired.

We bound $\Pr[X_{km}(e) \neq Y_{km}(e)]$ for the case where $X_0 = E_n$ and Y_0 is sampled from μ^η. The proof for the case where $X_0 = \emptyset$ and Y_0 is sampled from μ^η is identical. The monotonicity of the identity coupling discussed in Section 2 implies then that this bound holds for arbitrary initial configurations X_0, Y_0.

Let $\mathcal{E}_k$ be the event $\{X_{km}(B(e,r)) = Y_{km}(B(e,r))\}$ for some r we will choose later, and let

$$\rho(k) := \max_{e \in E_n} \Pr[\, X_{km}(e) \neq Y_{km}(e)\,],$$

where the probability is over both the choice of Y_0 and the steps of the Markov chain. We will show that $\rho(k) \leq \exp(-\Omega(k))$.

Since Y_0 has law μ^η_Λ, the bound on disagreement percolation in Lemma 3.1 implies that

$$\Pr[\, X_{2km}(e) \neq Y_{2km}(e) \,|\, \mathcal{E}_k\,] \leq \mathrm{e}^{-\Omega(r^{1/4})},$$

provided $1 \leq k \leq r^{1/4}/(4\mathrm{e}^2)$ and r is large enough. Thus,

$$\Pr[\, X_{2km}(e) \neq Y_{2km}(e)\,] \leq \Pr[\, X_{2km}(e) \neq Y_{2km}(e) \,|\, \neg\mathcal{E}_k\,] \Pr[\, \neg\mathcal{E}_k\,] + \mathrm{e}^{-\Omega(r^{1/4})}.$$

The monotonicity of the identity coupling implies that $\Pr[\, X_{2km}(e) \neq Y_{2km}(e) \,|\, \neg\mathcal{E}_k\,] \leq \rho(k)$, and a union bound over the edges in $B(e,r)$ implies $\Pr[\, \neg\mathcal{E}_k\,] \leq \Theta(r^2)\rho(k)$. Putting these bounds together and taking $r = \Theta(k^4)$, we obtain

$$\rho(2k) \leq Ck^8\rho^2(k) + \mathrm{e}^{-\lambda k},$$

for some suitable constants $C > 1$ and $\lambda > 0$. (Note that since $r = \Theta(k^4)$ and $r < n$, our proof of inequality (13) does not hold for arbitrarily large k; hence the restriction $k = o(m^{1/8})$.)

Now, let

$$\phi(k) := 2^8(Ck^8 + 1)\max\{\rho(k), \mathrm{e}^{-\lambda k/2}\}.$$

We show next that $\phi(2k) \leq \phi(k)^2$. For this observe that $2^8(C(2k)^8 + 1)\rho(2k) \leq \phi(k^2)$. Hence, if $\rho(2k) \geq e^{-\lambda k}$, we get $\phi(2k) \leq \phi(k^2)$ directly. Otherwise, we have

$$\phi(2k) \leq (2^8)^2(Ck^8 + 1)\mathrm{e}^{-\lambda k} \leq \phi(k)^2.$$

Thus, for any integer $\alpha > 0$, we get $\phi(k) \leq \phi(k/2^\alpha)^{2^\alpha}$. The result follows from the following fact, which provides a stopping point for this recurrence.

Claim 5.3. *Let $l = m_0^{1/8}/A$, with $A = 2(8C\mathrm{e})^{1/8}$. Then $\rho(l) \leq \frac{1}{2^8\,\mathrm{e}(Cl^8+1)}$ for a sufficiently large m_0.*

As a result, $\phi(l) \leq 1/\mathrm{e}$ for a sufficiently large constant l, and thus $\rho(k) \leq \phi(k) \leq \mathrm{e}^{-k/l}$ as desired. □

We conclude this section with the proof of Claim 5.3. The proof is similar to that of Theorem 5.1 and makes crucial use of the hypothesis on the mixing time in square boxes of volume m_0.

Proof. Let $e \in E_n$ and choose r' such that $|E(e,r')| = m_0$. (Note that as a result $r' = \Theta(m_0^{1/2})$.) The proof proceeds along the same lines as that of Theorem 5.1. In fact we consider the same processes $\{Z_t^+\}$, $\{Z_t^-\}$, where $Z_0^+ = X_0 = E_n$, $Z_t^- = \emptyset$ and $\{Z_t^+\}$, $\{Z_t^-\}$ only update edges with both endpoints in $B(e,r')$. As before, we couple the four chains $\{X_t\}$, $\{Y_t\}$, $\{Z_t^+\}$, $\{Z_t^-\}$ with

the identity coupling, ignoring the updates outside $B(e, r')$ in $\{Z_t^+\}$ and $\{Z_t^-\}$. The monotonicity of the identity coupling then implies that $Z_t^- \subseteq Y_t \subseteq X_t \subseteq Z_t^+$ for all $t \geq 0$. Hence, we obtain inequality (10)-(12). (Note that in this case Y_t has law μ_Λ^η for all $t \geq 0$.)

Lemma 4.1 implies that (11) is at most $\exp(-\Omega(r'))$. To bound (10), note that if we run the identity coupling for lm steps, a Chernoff bound implies that with probability at least $1 - \exp(-lm_0/8)$ the number of updates in $B(e, r')$ is at least $(2A)^{-1}m_0^{9/8}$. If indeed this many steps hit $B(e, r')$, then the hypothesis of the theorem (with $c = (2A)^{-1}$) implies

$$\big|\Pr[\, Z_t^+(e) = 1\,] - \mu_\Lambda^\eta(\, e = 1 \,|\, E^c = 1)\big| \leq ||\, Z_t^+(\cdot) - \mu_\Lambda^\eta(\,\cdot \,|\, E^c = 1\,)\,||_{\mathrm{TV}} \leq \frac{1}{m_0}.$$

(Here we also used the fact that $\tau_{\mathrm{mix}}(\varepsilon) \leq \lceil \log_2 \varepsilon^{-1} \rceil \tau_{\mathrm{mix}}$.) We do the same to bound (12), and then

$$\Pr[\, X_{lm}(e) \neq Y_{lm}(e)\,] \leq \frac{2}{m_0} + e^{-\Omega(m_0^{9/8})} + e^{-\Omega(m_0^{1/2})} \leq \frac{4}{m_0} = \frac{4}{(Al)^8} \leq \frac{1}{2^8\,\mathrm{e}(Cl^8+1)},$$

for a sufficiently large constant m_0. Hence, $\phi(l) \leq 1/\,\mathrm{e}$ for a sufficiently large l, as desired. □

6 Mixing time lower bound in the sub-critical regime

In this section we prove the lower bound from Theorem 1.1 for $p < p_c(q)$. (The lower bound for $p > p_c(q)$ is derived in Section 7.) In the setting of spin systems, [16] provides a general mixing time lower bound for Glauber dynamics. As mentioned earlier, the random-cluster model is not a spin system in the usual sense because of the long range interactions, but we are still able to adapt the techniques in [16] to the random-cluster setting. In fact, our proof follows closely the argument in the proof of Theorem 4.1 in [16], the main difference being that we require a more subtle choice of the starting configuration to limit the effect of the long range interactions. We derive the following theorem.

Theorem 6.1. *Let $p < p_c(q)$, $q \geq 1$ and let η be a side-homogeneous boundary condition for $\Lambda = (\Lambda_n, E_n)$. The mixing time of the Glauber dynamics in Λ is $\Omega(n^2 \log n)$.*

It is convenient to carry out our proof in continuous time. The continuous time Glauber dynamics is obtained by adding a rate 1 Poisson clock to each edge; when the clock at edge e rings, e is updated as in discrete time.

The switch to continuous time requires us to extend the bound in Section 3 for the speed of propagation of disagreements to the continuous time dynamics. In addition, we will require slightly different assumptions about the initial configuration Y_0. This is established in the following lemma, whose proof is very similar to that of Lemma 3.1 (only requiring minor adjustments) and is deferred to Appendix A.

Lemma 6.2. *Let $p < p_c(q)$, $q \geq 1$ and let η be a side-homogeneous boundary condition for $\Lambda = (\Lambda_n, E_n)$. Also, let $B = B(e, r)$ for some $e \in E_n$ and $r \geq 1$. Consider two copies $\{X_t\}$, $\{Y_t\}$ of the continuous time Glauber dynamics in Λ such that:*

- *$X_0(B) = Y_0(B)$;*
- *$Y_0(B)$ has law $\mu^0_{B,p,q}(\,\cdot \,|\, e = b)$ for some $b \in \{0, 1\}$;*
- *$Y_0(e') = 0$ for all $e' \in E^c(e, r)$ incident to ∂B; and*

- $\{Y_t\}$ *only performs edge updates in* B.

If the evolutions of $\{X_t\}$ *and* $\{Y_t\}$ *are coupled using the identity coupling, then there exist positive constants* c, C *and* λ *independent of* r *such that, for all* $r \geq c$ *and* $1 \leq T \leq r^{1/4}/(4\mathrm{e}^2)$, *we have*

$$\Pr[\, X_T(e) \neq Y_T(e) \,] \;\leq\; C\mathrm{e}^{-\lambda r^{1/4}}.$$

We are now ready to prove Theorem 6.1.

Proof. Let $\{X_t\}$ and $\{X_t^{\mathcal{D}}\}$ be copies of the continuous and discrete time Glauber dynamics in Λ, respectively, such that $X_0^{\mathcal{D}} = X_0$. The following standard inequality holds for all $t \geq 0$:

$$||X_{t'}^{\mathcal{D}} - \mu^\eta||_{\mathrm{TV}} \geq ||X_t - \mu^\eta||_{\mathrm{TV}} - 2e^{-t'},$$

where $t' = |E_n|t/2$ (see, e.g., Proposition 2.1 in [16]).

We will show that $||X_T - \mu^\eta||_{\mathrm{TV}} > 1/3$ for some $T = \Omega(\log n)$; as a result $||X_{T'}^{\mathcal{D}} - \mu^\eta||_{\mathrm{TV}} > 1/4$ for some $T' = \Omega(n^2 \log n)$ and sufficiently large n. This implies that the mixing time of the discrete time dynamics is $\Omega(n^2 \log n)$ as desired.

First we introduce some notation. Assume w.l.o.g. that $4r+1$ divides $n-1$ for some r to be chosen later, and split Λ_n into $(n-1)^2/(4r+1)^2$ square boxes of side length $4r+1$. Each of these boxes corresponds to $B(e, 2r)$ for some edge $e \in E_n$; let $\mathcal{C} \subseteq E_n$ be the set of these edges. Also, let $\hat{E} := E_n \setminus \bigcup_{e \in \mathcal{C}} E(e, r)$ and let $\mathcal{E}$ be the event that every edge $e' \in \hat{E}$ incident to $B(e, r)$ for some $e \in \mathcal{C}$ is closed.

Let A, $A_{\mathcal{E}}$ be random-cluster configurations sampled from μ^η and $\mu^\eta(\,\cdot \,|\, \mathcal{E})$, respectively, and let $\beta := \mathrm{E}[f(A_{\mathcal{E}})]$, where $f(A_{\mathcal{E}})$ is the fraction of edges $e \in \mathcal{C}$ such that $A_{\mathcal{E}}(e) = b$ for some fixed $b \in \{0, 1\}$. Consider the following threshold for a value of $\varepsilon > 0$ that will be chosen later:

$$\hat{\beta} = \begin{cases} \beta + \varepsilon & \text{if } \beta < 1/2; \\ \beta & \text{if } \beta = 1/2; \\ \beta - \varepsilon & \text{if } \beta > 1/2. \end{cases}$$

We pick b such that $\Pr[f(A) > \hat{\beta}] \leq 1/2$.

As in [16], our goal is to choose X_0 such that $\Pr[f(X_T) \geq \hat{\beta}] \to 1$ as $n \to \infty$ for some $T = \Omega(\log n)$; then $||X_T - \mu^\eta||_{\mathrm{TV}} > 1/3$ for large enough n, as desired.

Let Φ be the set of random-cluster configurations in Λ such that $\hat{A} \in \Phi$ iff for all $e \in \mathcal{C}$, $\hat{A}(e) = b$ and $\hat{A}(e') = 0$ for all $e' \in \hat{E}$ incident to $B(e, r)$. For each $\hat{A} \in \Phi$, let

$$\pi_0(\hat{A}) := \frac{\mu^\eta(\hat{A})}{\mu^\eta(\Phi)}. \tag{14}$$

The starting configuration X_0 is sampled from π_0.

Consider now an auxiliary copy $\{Y_t\}$ of the continuous time Glauber dynamics such that $Y_0 = X_0$. The two chains $\{X_t\}$,$\{Y_t\}$ are coupled using the identity coupling, except that $\{Y_t\}$ does not update edges in $\hat{E}$. First we establish a bound for $\Pr[f(Y_T) \leq \hat{\beta} + \varepsilon/2]$. To do this, we use the following monotonicity property which is a straightforward consequence of Lemma 3.5 in [16].

Fact 6.3. *For each* $e \in \mathcal{C}$, *let* $\alpha_e := \mu^0_{B,p,q}(e = b)$ *where* $B = B(e, r)$. *Then, for all* $t \geq 0$,

$$\Pr[Y_t(e) = b] \geq \alpha_e + (1 - \alpha_e)\mathrm{e}^{-t/(1-\alpha_e)}.$$

From this fact, we follow the steps in the proof of Theorem 4.1 in [16] to derive the following bound:

$$\mathrm{E}[f(Y_T)] \geq \beta + (1-\beta)\mathrm{e}^{-T/(1-\beta)}, \tag{15}$$

for all $T \geq 0$. Taking $\varepsilon = 1/(4\exp(2T))$, the right hand side of (15) is at least $\hat{\beta}+\varepsilon$. Also, since Y_0 is sampled from π_0, the configurations of the edges in $\mathcal{C}$ are independent in Y_T. Hence, $|\mathcal{C}|f(Y_T)$ is the sum of $|\mathcal{C}|$ independent Bernoulli random variables; a Chernoff bound then implies

$$\Pr[f(Y_T) \leq \hat{\beta}+\varepsilon/2] \leq \mathrm{e}^{-\Omega(\varepsilon^2|\mathcal{C}|)}. \tag{16}$$

The second step in the proof is to bound $\Pr[f(X_T) \leq f(Y_T) - \varepsilon/2]$. For this, we use Lemma 6.2, which is tailored precisely to our setting. Thus, for all $e \in \mathcal{C}$ and $1 \leq T \leq r^{1/4}/(4\mathrm{e}^2)$, we have

$$\Pr[X_T(e) \neq Y_T(e)] \leq C\mathrm{e}^{-\lambda r^{1/4}},$$

provided $r \geq c$, where c is a sufficiently large constant. Consequently, the expected number of disagreements between X_T and Y_T in $\mathcal{C}$ is at most $|\mathcal{C}|C\mathrm{e}^{-\lambda r^{1/4}}$, and by Markov's inequality,

$$\Pr[f(X_T) \leq f(Y_T) - \varepsilon/2] \leq \frac{2C}{\varepsilon}\mathrm{e}^{-\lambda r^{1/4}}. \tag{17}$$

Putting together the bounds in (16) and (17), we get

$$\Pr[f(X_T) \leq \hat{\beta}] \leq \frac{2C}{\varepsilon}\mathrm{e}^{-\lambda r^{1/4}} + \mathrm{e}^{-\Omega(\varepsilon^2|\mathcal{C}|)}.$$

Finally, observe that $|\mathcal{C}| = \Theta(\frac{n^2}{r^2})$; thus, when $r = (\frac{1}{4}\log n)^4$ and $T = \min\{\lambda/4, 1\}r^{1/4}$, we get $\Pr[f(X_T) > \hat{\beta}] \to 1$ as $n \to \infty$ as desired. □

7 Mixing time in the super-critical regime

In this section we prove Theorem 1.1 from the Introduction for $p > p_c(q)$. We will make use of planar duality, discussed in Section 2, in order to reduce the proof to the sub-critical case.

Theorem 7.1. *Let $p > p_c(q)$ and $q \geq 1$. Then the mixing time of the Glauber dynamics on $\Lambda = (\Lambda_n, E_n)$, with free or wired boundary conditions, is $\Theta(n^2 \log n)$.*

Proof. We focus on the free boundary condition case; the wired case follows from an analogous argument. The planar dual $\Lambda^* = (\Lambda_n^*, E_n^*)$ of the graph Λ consists of an $(n-1)\times(n-1)$ box with an additional outer vertex corresponding to the infinite face of Λ. The dual measure $\mu_{\Lambda^*,p^*,q}$, with

$$p^* = \frac{q(1-p)}{p+q(1-p)},$$

is equivalent to the measure $\mu^1_{\Lambda',p^*,q}$ where $\Lambda' = (\Lambda_{n+1}, E_{n+1})$ (see, e.g., [3]). Note that $p > p_c(q)$ iff $p^* < p_c(q)$.

We say that two random-cluster configurations A on Λ and A' on Λ' are *compatible* if the configuration resulting from A' by contracting all the vertices in the boundary of Λ' into a single vertex is A^*, the dual configuration of A. Note that each A' has a unique compatible A, while each A has multiple compatible A' that differ only in the disposition of the edges in the boundary $\partial\Lambda_{n+1}$. Observe also that any edge e' of Λ' with at most one endpoint incident to $\partial\Lambda_{n+1}$ corresponds to a unique dual edge $e^* \in E_n^*$ and thus to a unique edge $e \in E_n$.

In order to analyze the Glauber dynamics on Λ when $p > p_c(q)$, we consider instead the Glauber dynamics on Λ' with parameter $p^* < p_c(q)$, which we denote $\{X'_t\}$. We shall show that $\{X'_t\}$ induces a Markov chain $\{X_t\}$ on Λ which is essentially the same as the Glauber dynamics on Λ with parameter p, and that the mixing times of $\{X_t\}$ and $\{X'_t\}$ are equal up to constant factors. Since $p^* < p_c(q)$, the results in Sections 5 and 6 imply that mixing time $\{X'_t\}$ (and hence of $\{X_t\}$) is $\Theta(n^2 \log n)$.

To define the induced dynamics, let e'_t be the edge chosen u.a.r. from E_{n+1} at time t by $\{X'_t\}$, and let e_t be the corresponding edge in Λ if there is one. X_{t+1} is obtained from X_t as follows:

(i) if both endpoints of e'_t are in $\partial\Lambda_{n+1}$, then $X_{t+1} = X_t$;

(ii) else if $X'_{t+1} = X'_{t+1} \cup \{e'_t\}$, then $X_{t+1} = X_t \setminus \{e_t\}$;

(iii) else if $X'_{t+1} = X'_{t+1} \setminus \{e'_t\}$, then $X_{t+1} = X_t \cup \{e_t\}$.

The initial configuration X_0 is the unique configuration compatible with X'_0.

We show first that $\{X_t\}$ is in fact a lazy version of the Glauber dynamics on Λ. To see this, note that $X_{t+1} = X_t$ whenever both endpoints of e'_t are in $\partial\Lambda_{n+1}$. Otherwise, it is straightforward to check that $e_t \in X_t$ is a cut edge iff $e'_t \in X'_t$ is not a cut edge. Hence, $X_{t+1} = X_t \cup \{e_t\}$ iff $X'_{t+1} = X_t \setminus \{e'_t\}$ and thus $X_{t+1} = X_t \cup \{e_t\}$ with probability:

$$\begin{cases} 1 - \frac{p^*}{q(1-p^*)+p^*} = p & \text{if } e'_t \text{ is a cut edge;} \\ 1 - p^* = \frac{p}{q(1-p)+p} & \text{otherwise.} \end{cases}$$

This implies that $\{X_t\}$ does not move with probability $\Theta(n^{-1})$, and otherwise evolves exactly like the Glauber dynamics on Λ. Hence, it is sufficient for us to establish the mixing time of $\{X_t\}$. To do this, we show that the mixing times of $\{X_t\}$ and $\{X'_t\}$ are essentially the same.

Let Ω_{RC} be the set of random-cluster configurations on Λ, and let $\mathcal{C}(A)$ be the set of configurations compatible with a configuration A on Λ. The first observation is that when $\{X'_t\}$ mixes, so does $\{X_t\}$. This follows from:

$$\begin{aligned} ||\, X_t(\cdot) - \mu_{\Lambda,p,q}(\cdot)\,||_{\mathrm{TV}} &= \frac{1}{2} \sum_{A \in \Omega_{\mathrm{RC}}} |X_t(A) - \mu_{\Lambda,p,q}(A)| \\ &= \frac{1}{2} \sum_{A \in \Omega_{\mathrm{RC}}} |X_t(A) - \mu_{\Lambda^*,p^*,q}(A^*)| \\ &\leq \frac{1}{2} \sum_{A \in \Omega_{\mathrm{RC}}} \sum_{A' \in \mathcal{C}(A)} \left| X'_t(A') - \mu^1_{\Lambda',p^*,q}(A') \right| \\ &= ||\, X'_t(\cdot) - \mu^1_{\Lambda',p^*,q}(\cdot)\,||_{\mathrm{TV}}, \end{aligned}$$

where in the first and last equality we use the definition of total variation distance, in the second equality we use planar duality, and the third inequality follows from the triangle inequality and the correspondence between the configurations of Λ and Λ'. Hence, by the results in Section 5, the mixing time of $\{X_t\}$ is $O(n^2 \log n)$.

We show next that the mixing time of $\{X_t\}$ is $\Omega(n^2 \log n)$. For this, note that in Theorem 6.1 we showed that there is an initial distribution π_0 for X'_0, defined in (14), such that

$$||X'_T(\cdot) - \mu^1_{\Lambda',p^*,q}(\cdot)||_{\mathrm{TV}} > 1/4$$

for some $T = \Omega(n^2 \log n)$. We will prove that when X'_0 is sampled from π_0, then

$$||X'_t(\cdot) - \mu^1_{\Lambda',p^*,q}(\cdot)||_{\mathrm{TV}} = ||X_t(\cdot) - \mu_{\Lambda,p,q}(\cdot)||_{\mathrm{TV}}$$

for all $t \geq 0$. To show this we introduce some additional notation.

Let $\Lambda'' := (\Lambda_{n+1}, E_{n+1} \setminus \partial E_{n+1})$, where $\partial E_{n+1} \subseteq E_{n+1}$ is the set of edges with both endpoints in $\partial\Lambda_{n+1}$. Also, for any random-cluster configuration A' on Λ', we use $\partial A'$ to denote the random-cluster configuration induced in $\partial\Lambda'$ by A'.

Under the wired boundary condition, we have that for any $e \in \partial\Lambda'$, $\mu^1_{\Lambda',p^*,q}(e = 1) = p^*$. Hence, $\mu^1_{\Lambda',p^*,q}$ is the product measure of the distributions $\mu^1_{\Lambda'',p^*,q}$ and $\mu_{\partial\Lambda',p^*,1}$; the latter is the distribution on $\partial\Lambda'$ where every edge is sampled independently with probability p^*. Thus we have

$$\mu^1_{\Lambda',p^*,q}(A') = \mu^1_{\Lambda'',p^*,q}(A' \setminus \partial A') \cdot \mu_{\partial\Lambda',p^*,1}(\partial A').$$

By the correspondence between the configurations of Λ^* and Λ', $\mu^1_{\Lambda'',p^*,q}(A' \setminus \partial A') = \mu_{\Lambda^*,p^*,q}(A^*)$. (As in Section 2, A^* denotes the dual of the unique configuration A compatible with A'.) Moreover, by planar duality $\mu_{\Lambda^*,p^*,q}(A^*) = \mu_{\Lambda,p,q}(A)$, and thus

$$\mu^1_{\Lambda',p^*,q}(A') = \mu_{\Lambda,p,q}(A) \cdot \mu_{\partial\Lambda',p^*,1}(\partial A'). \tag{18}$$

Also, under the wired boundary condition the configuration on the boundary of X'_0 is sampled according to $\mu_{\partial\Lambda',p^*,1}$. Hence, the distribution on the boundary of X'_t has law $\mu_{\partial\Lambda',p^*,1}$ for all $t \geq 0$. Thus,

$$X'_t(A') = \mu_{\partial\Lambda',p^*,1}(\partial A') \cdot X'_t(A' \setminus \partial A') = \mu_{\partial\Lambda',p^*,1}(\partial A') \cdot X_t(A). \tag{19}$$

Hence,

$$\begin{aligned} || X'_t(\cdot) - \mu^1_{\Lambda',p^*,q}(\cdot) ||_{\mathrm{TV}} &= \frac{1}{2} \sum_{A \in \Omega_{\mathrm{RC}}} \sum_{A' \in \mathcal{C}(A)} \left| X'_t(A') - \mu^1_{\Lambda',p^*,q}(A') \right| \\ &= \frac{1}{2} \sum_{A \in \Omega_{\mathrm{RC}}} \sum_{A' \in \mathcal{C}(A)} \mu_{\partial\Lambda',p^*,1}(\partial A') \, |X_t(A) - \mu_{\Lambda,p,q}(A)| \\ &= || X_t(\cdot) - \mu_{\Lambda,p,q}(\cdot) ||_{\mathrm{TV}}, \end{aligned}$$

where in the first and last equality we used the definition of total variation distance and the second follows from (18) and (19).

The results in Section 6 then imply that the mixing time of $\{X_t\}$ is $\Omega(n^2 \log n)$. Consequently, the Glauber dynamics on Λ with $p > p_c(q)$ mixes in $\Theta(n^2 \log n)$ steps, as desired. □

Acknowledgments

The authors would like to thank Fabio Martinelli and Allan Sly for helpful suggestions. We would like also to thank an anonymous referee whose valuable comments improved the exposition of the results.

References

[1] K.S. Alexander. On weak mixing in lattice models. *Probability Theory and Related Fields*, 110:441–471, 1998.

[2] K.S. Alexander. Mixing properties and exponential decay for lattice systems in finite volumes. *The Annals of Probability*, 32(1A):441–487, 2004.

[3] V. Beffara and H. Duminil-Copin. The self-dual point of the two-dimensional random-cluster model is critical for $q \geq 1$. *Probability Theory and Related Fields*, 153:511–542, 2012.

[4] J. van den Berg. A uniqueness condition for Gibbs measures, with applications to the 2-dimensional ising antiferromagnet. *Communications in Mathematical Physics*, 152(1):161–166, 1993.

[5] A. Blanca and A. Sinclair. Dynamics for the mean-field random-cluster model. *Proceedings of the 19th International Workshop on Randomization and Computation*, pages 528–543, 2015.

[6] C. Borgs, J. Chayes, and P. Tetali. Swendsen-Wang algorithm at the Potts transition point. *Probability Theory and Related Fields*, 152:509–557, 2012.

[7] C. Borgs, A. Frieze, J.H. Kim, P. Tetali, E. Vigoda, and V. Vu. Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics. *Proceedings of the 40th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 218–229, 1999.

[8] H. Duminil-Copin, V. Sidoravicius, and V. Tassion. Continuity of the phase transition for planar random-cluster and Potts models with $1 \leq q \leq 4$. *Preprint*, 2015.

[9] M. Dyer, A. Sinclair, E. Vigoda, and D. Weitz. Mixing in time and space for lattice spin systems: A combinatorial view. *Random Structures & Algorithms*, 24:461–479, 2004.

[10] R.G. Edwards and A.D. Sokal. Generalization of the Fortuin-Kasteleyn-Swendsen-Wang representation and Monte Carlo algorithm. *Physical Review D*, 38(6):2009–2012, 1988.

[11] C.M. Fortuin. On the random-cluster model: III. The simple random-cluster model. *Physica*, 59(4):545–570, 1972.

[12] C.M. Fortuin and P.W. Kasteleyn. On the random-cluster model I. Introduction and relation to other models. *Physica*, 57(4):536–564, 1972.

[13] A. Galanis, D. Štefankovič, and E. Vigoda. Swendsen-Wang algorithm on the mean-field Potts model. *Proceedings of the 19th International Workshop on Randomization and Computation*, pages 815–828, 2015.

[14] Q. Ge and D. Štefankovič. A graph polynomial for independent sets of bipartite graphs. *Combinatorics, Probability and Computing*, 21(5):695–714, 2012.

[15] G.R. Grimmett. *The Random-Cluster Model*, volume 333 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 2006.

[16] T.P. Hayes and A. Sinclair. A general lower bound for mixing of single-site dynamics on graphs. *Annals of Applied Probability*, 17(3):931–952, 2007.

[17] L. Laanait, A. Messager, S. Miracle-Solé, J. Ruiz, and S. Shlosman. Interfaces in the Potts model I: Pirogov-Sinai theory of the Fortuin-Kasteleyn representation. *Communications in Mathematical Physics*, 140(1):81–91, 1991.

[18] D.A. Levin, Y. Peres, and E.L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

[19] Y. Long, A. Nachmias, W. Ning, and Y. Peres. A power law of order 1/4 for critical mean-field Swendsen-Wang dynamics. *Memoirs of the American Mathematical Society*, 232(1092), 2011.

[20] E. Lubetzky and A. Sly. Critical ising on the square lattice mixes in polynomial time. *Communications in Mathematical Physics*, 313(3):815–836, 2012.

[21] F. Martinelli. *Lectures on Glauber dynamics for discrete spin models*, volume 1717 of *Springer Lecture Notes in Mathematics*. Springer Verlag, 1999.

[22] F. Martinelli, E. Olivieri, and R.H. Schonmann. For 2-d lattice spin systems weak mixing implies strong mixing. *Communications in Mathematical Physics*, 165(1):33–47, 1994.

[23] E. Mossel and A. Sly. Exact thresholds for Ising–Gibbs samplers on general graphs. *The Annals of Probability*, 41(1):294–328, 2013.

[24] J. Propp and D. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures & Algorithms*, 9:223–252, 1996.

[25] V. Strassen. The existence of probability measures with given marginals. *The Annals of Mathematical Statistics*, 36:423–439, 1965.

[26] R.H. Swendsen and J.S. Wang. Nonuniversal critical dynamics in Monte Carlo simulations. *Physical Review Letters*, 58:86–88, 1987.

[27] M. Ullrich. Comparison of Swendsen-Wang and heat-bath dynamics. *Random Structures and Algorithms*, 42(4):520–535, 2013.

[28] M. Ullrich. Rapid mixing of Swendsen-Wang dynamics in two dimensions. *Dissertationes Mathematicae*, 502, 2014.

[29] M. Ullrich. Swendsen-Wang is faster than single-bond dynamics. *SIAM Journal on Discrete Mathematics*, 28(1):37–48, 2014.

A Proof of Lemma 6.2

We show first that the measure that results from conditioning on the state of a single edge maintains the exponential decay of finite volume connectivities (3).

Fact A.1. *Let $p < p_c(q)$, $q \geq 1$, and let η be a boundary condition for $\Lambda = (\Lambda_n, E_n)$. Consider a copy $\{Y_t\}$ of the continuous time Glauber dynamics on Λ, and assume Y_0 is sampled from the distribution $\mu^{\eta}_{\Lambda,p,q}(\,\cdot \,|\, e = b)$, for some $e \in E_n$ and $b \in \{0,1\}$. Then, for all $u, v \in \Lambda_n$, there exists positive constant C and λ such that*

$$\Pr[u \overset{Y_t}{\leftrightarrow} v] \leq C\mathrm{e}^{-\lambda\, d(u,v)},$$

where $u \overset{Y_t}{\leftrightarrow} v$ denotes the event that u and v are connected by a path of open edges in Y_t.

Proof. Let $\{Z_t\}$ be a second instance of the continuous time Glauber dynamics. The evolution of $\{Z_t\}$ is coupled with that of $\{Y_t\}$ via the identity coupling, except that $\{Z_t\}$ never updates the edge e. The initial configuration of $\{Z_t\}$ is sampled according to the distribution $\mu^\eta(\,\cdot\,|\,e=1)$ such that $Y_0 \subseteq Z_0$. This is always possible because $\mu^\eta(\,\cdot\,|\,e=1) \succeq \mu^\eta$. Then, $Y_t \subseteq Z_t$ and Z_t has law $\mu^\eta(\,\cdot\,|\,e=1)$ for all $t \geq 0$. We establish that the measure $\mu^\eta(\,\cdot\,|\,e=1)$ has exponential decay of finite volume connectivities and thus so does the distribution of Y_t for all $t \geq 0$. By (3), for all $u, v \in \Lambda_n$, we have

$$\mu^\eta(\,u \leftrightarrow v\,|\,e=1)\mu^\eta(e=1) \leq \mu^\eta(u \leftrightarrow v) \leq C\mathrm{e}^{-\,\lambda\, d(u,v)},$$

where C, λ are positive constants. If $p' = \frac{p}{q(1-p)+p}$, then $\mu^\eta \succeq \mu^\eta_{\Lambda,p',1}$ (see, e.g., [11]), and thus $\mu^\eta(e=1) \geq p'$. Since $q \geq 1$,

$$\mu^\eta(\,u \leftrightarrow v\,|\,e=1) \leq \frac{qC}{p}\mathrm{e}^{-\,\lambda\, d(u,v)}.$$

The result then follows immediately when $p = \Omega(1)$. Otherwise, the measure μ^η is stochastically dominated by any random-cluster measure $\mu^\eta_{\Lambda,p'',q}$ with $p'' = \Omega(1)$, for which we just established exponential decay of finite volume connectivities; the result follows by monotonicity. □

We are now ready to prove the lemma.

Proof. Let Q_t be the random time at which the t-th edge is updated by the identity coupling. For some fixed ℓ to be chosen later, and each $t \geq 0$, consider the event

$$\mathcal{E}_{\ell,t} := \{u \overset{Y_{Q_t}}{\nleftrightarrow} v \quad \forall u, v \in B \text{ s.t. } d(u,v) > \ell\},$$

where $u \overset{Y_{Q_t}}{\nleftrightarrow} v$ denotes the event that u and v are not connected by a path in $Y_{Q_t}(B)$. Also, let $\mathcal{E}_\ell := \bigcap_{t:Q_t \leq T} \mathcal{E}_{\ell,t}$. Then,

$$\Pr[\,X_T(e) \neq Y_T(e)\,] \leq \Pr[\,X_T(e) \neq Y_T(e)\,|\,\mathcal{E}_\ell\,] + \Pr[\,\neg\mathcal{E}_\ell\,] \tag{20}$$

(cf. equation (5)). We bound each term on the right hand side of (20) separately.

Conditioned on the event $\mathcal{E}_\ell$, a witness for the fact that $X_T(e) \neq Y_T(e)$ can be constructed as in discrete time. However, the probability that a given witness of length L is updated by the continuous time dynamics is instead bounded using the following fact from [16].

Fact A.2. *Consider L independent rate 1 Poisson clocks. Then, the probability that there is an increasing sequence of times $0 \leq t_1 < \ldots < t_L \leq T$ such that clock i rings at time t_i is at most $\left(\frac{eT}{L}\right)^L$.*

Recall from Section 3 that the number of witnesses of length L is (crudely) at most $(4(\ell+1)^2)^L$. Hence, following the same steps as in the proof Lemma 3.1, and taking $\ell = r^{1/4} - 1$, we get

$$\Pr[\,X_T(e) \neq Y_T(e)\,|\,\mathcal{E}_\ell\,] \leq \frac{\mathrm{e}}{\mathrm{e}-1} \cdot \mathrm{e}^{-r^{3/4}}, \tag{21}$$

using the fact that $T \leq r^{1/4}/(4\mathrm{e}^2)$ (cf. equation (7)).

To bound the second term on the right hand side of (20), let N be the number of edge updates in B up to time T. Observe that N is a Poisson random variable with rate $M := |E(B,r)| = \Theta(r^2)$. Using standard bounds for Poisson tail probabilities we get $\Pr[N > \mathrm{e}^2 MT] = \exp(-\Omega(MT))$ for all $T \geq 1$. Therefore,

$$\Pr[\,\neg\mathcal{E}_\ell\,] \leq \Pr[\,\neg\mathcal{E}_\ell\,|\,N \leq \mathrm{e}^2 MT\,] + \mathrm{e}^{-\Omega(MT)}.$$

Also, $\neg\mathcal{E}_\ell := \bigcup_{t:Q_t \le T} \neg\mathcal{E}_{\ell,t}$, and if the edge update at time Q_t occurs outside B, we have $\neg\mathcal{E}_{\ell,t} = \neg\mathcal{E}_{\ell,t+1}$. Hence, a union bound implies

$$\Pr[\,\neg\mathcal{E}_\ell\,] \;\le\; \mathrm{e}^2 MT \max_{t:Q_t \le T} \Pr[\,\neg\mathcal{E}_{\ell,t}\,] + \mathrm{e}^{-\Omega(MT)}.$$

Fact A.1 establishes exponential decay of finite volume connectivities (3) for the distribution of Y_t in B for all $t \ge 0$. Then, as in Lemma 3.1, we obtain

$$\Pr[\,\neg\mathcal{E}_\ell\,] \le O(r^{6.25}) \cdot \mathrm{e}^{-\Omega(r^{1/4})} + \mathrm{e}^{-\Omega(r^2)}.$$

Together with (21), this implies there exist constants $c, C, \lambda > 0$ such that for all $r \ge c$ we have $\Pr[\,X_T(e) \ne Y_T(e)\,] \le C\mathrm{e}^{-\lambda r^{1/4}}$, as desired. □